Intelligent RF-IR Data Fusion for Space Objects using Artificial Intelligence Techniques

Richard Stottler[1], Edward Watson, and Michael Wicks
Stottler Henke Associates, Inc., San Mateo, CA and University of Dayton Research Institute, Dayton, OH

We developed techniques and software to automatically extract high-level object features (such as object rotational frequency and size) from Infrared (IR) and radar signals, specifically when the objects are less than the size of a single IR pixel or radar beam, and executed them on actual radar and IR sensor data. These features can be used to automatically and intelligently correlate the same object in different scenes sensed by dissimilar sensors and the additional data used to make better classification decisions. These techniques were implemented in an automatic-reasoning prototype, which could autonomously extract features from the IR and radar signals, use multiple features to correlate objects in different sensor scenes, calculate probabilities that two different sensed objects in different scenes were, in fact, the same object, use these probabilities to construct the most probable hypothetical worlds, and extend these correlations and hypothetical worlds across any number of scenes, pruning the least probable worlds as required. There appears to be enormous potential for this work to have a major impact on correlation and classification of objects in space. Object rotational frequency, in particular, was shown to be extractable, very precisely and robustly providing substantial correlation power.

Nomenclature

DTFT = Discrete Time Fourier Transform
IR = Infrared
ISS = International Space Station
LWIR = Long Wave IR
MWIR = Midwave IR
SSN = Space Surveillance Network

I. Motivation

Tracking and classification of space objects is important for a number of reasons. The Air Force's Space Surveillance Network (SSN) consisting of optical, Infrared (IR), and radar sensors is used to track and classify nearly 20,000 man-made objects orbiting the Earth, called the Space Catalog. Accurately tracking and classifying these objects is important for the safety of objects in orbit such as functioning satellites and manned spacecraft. For example, the orbit of the International Space Station (ISS) is often adjusted based on the Space Catalog. There are large gaps in sensor coverage such that correlating objects in one scene with the same objects in another can be problematic, especially when the sensors are of different types. Similarly, for missile defense, which uses many of the same sensors as the SSN, being able to accurately track and classify objects is important to discriminate dangerous objects from nonlethal ones and provide accurate intercept data.

II. Overview

The effort involved two different thrusts. The first was experimental validation that object features such as size and rotational frequency could be accurately extracted from actual IR and radar signals from objects that were smaller than a single IR pixel and smaller than the radar's beam width. Experiments were performed with an actual radar and IR cameras simultaneously sensing various objects rotating at various speeds and angles. Features were extracted from each sensor signal and compared with an emphasis on rotational frequency since it was correctly anticipated that it would be the most precise.

[1] President, 1670 South Amphlett Blvd., Suite 310, San Mateo, CA, 94402, AIAA Member.

The second thrust was on developing techniques to intelligently and automatically correlate objects in different sensor scenes by reasoning over the extracted features. A software prototype was developed and executed on scenarios constructed using the data collected in the experiments. The prototype performed well. The automatic-reasoning prototype could autonomously extract features from the IR and radar signals, use multiple features to correlate objects in different sensor scenes, calculate probabilities that two different sensed objects in different scenes were, in fact, the same object, use these probabilities to construct the most probable hypothetical worlds, and extend these correlations and hypothetical worlds across any number of scenes, pruning the least probable worlds as required. Nine separate scenarios were constructed from the experiment runs to illustrate different capabilities of the prototype. Each scenario consisted of a set of observations and each observation referenced a set of IR or radar data files recorded from our experiments where each file represented the data from one object from one sensor for a specific condition (a specific experiment run). The set of scenarios ranged from simple examples (two scenes of two objects each) to ones that were more complex, including missing objects in some scenes, having 3 scenes to show that the prototype could handle any number, and including more objects, up to 8 in one case. The prototype successfully handled all situations presented to it, including use of multiple features, missing objects, more than 2 scenes, ambiguous situations, etc.

7 different shapes in two different states (unpainted and painted with high-emissivity paint, effectively creating 14 different targets) were used to gather radar and IR data for several experimental runs with a small variety of spin speeds and angles. Perhaps most remarkable was how robust the object rotation frequency extraction techniques were. We believe this is because of their simplicity. Good results were obtained for all conditions tested, even without calibrating the sensors or the use of past recorded data of any kind. All experiments used IR pixels and radar beam widths larger than the target objects. The average error in rotational frequency between IR and radar was an astounding 1.42% and much of this was attributable to the low-cost nature of the experiments and would not be present during operational use (e.g., actual variation in the rotational frequency of the object). Operational errors of our techniques are limited only by frame rate (more of a factor for IR than radar) and number of samples. Good results can be obtained with as little as 2 rotational periods (or even less in some Long Wave IR (LWIR) cases) of data. Other features besides rotational frequency were also extracted, including rotational size and center of gravity location from radar and average vertical and horizontal dimensions, and ratio of minimum and maximum subtended area from IR.

To our knowledge, this experiment represents the first time that there are simultaneous IR and radar signatures coupled with an inference algorithm to make object correlations and classification estimates. The real significance arises out of the fact that decoys or junk may have an RCS that looks like a target, or an IR signature that looks like a target, but probably not both at the same time. While the idea of using infrared measurements to record the time variation of rotating objects has been considered by the missile defense community before, their conclusion appears to have been that because the influence of projected area and material emissivity is coupled in the integrated irradiance signature that is measured, it is very difficult to make definitive statements about the object under observation. What we have shown is that heuristics can be used to derive useful features, even with little or no prior information on the objects or their sensor returns.

III. Experiments

In all experiments, objects were observed simultaneously by both IR cameras and the radar. This was to confirm that certain properties of the targets could be sensed and derived from a single pixel or single beam of data. In initial uncalibrated, raw form, the assumption was that the most obvious (and in some ways the most important) feature would be the rotational frequency of the object with the initial assumption that for most of our targets, the rotational frequency would be half the signal variation frequency. The full resolution of the IR cameras provided ground truth on the actual state of the targets. (I.e., it was straightforward to observe the rotational frequency of the targets on the IR video at full resolution.)

The initial plan was to send spinning targets down a zip line at night. Unfortunately, the windy conditions made the data difficult to interpret since the wind induced significant cross line motions (and were very unrealistic of actual conditions in space). The majority of the experiments were therefore performed nearer to the ground by using a large industrial fan that had had its blades removed and its input voltage lowered with a regulator so that its speed could be substantially decreased. Targets were attached to the fan's hub and spun at 2 or 3 different speeds and at 2 or 3 different angles each.

A. Stationary Rotations

A large industrial fan was on the premises, which had been used to test the radar as it was being constructed and the general principle of sensing oscillations with both radar and IR. It was determined that we could remove the fan blades and connect many of our targets to the hub. An attempt was made to have the object's center of gravity align with the spin axis, both to be the most realistic and to minimize wobble. Major advantages of this configuration over the zip line were:

- More constant rotational speed.
- Ability to easily vary the angle of rotation compared to the sensors.
- Far less wind-induced contamination of the signal.
- A more realistic space environment (where rotations should be nonvarying and, of course, there is no wind).
- Safer.

The major disadvantages were that the axis of rotation always would need to be horizontal, that the fan engine casing and hub were metal, which would create radar returns, and that there was more background clutter in the IR scene. Since neither the engine casing nor the background clutter was moving, these disadvantages could be handled by subtracting out the non-moving parts of the scene in the IR frames and ignoring the 0 Doppler returns in the radar signal (which was also required in the zip line experiments because the radar beam was wide enough that even when elevated up toward the zip line, side lobes returns from the earth caused a large 0 Doppler return). The 7 separate shapes were attached to rotate, one horizontal axis each, approximately through their center of gravity. Those 7 shapes were rotated and sensed both in their shiny metal (low-emissivity) surface state and when painted black with high-emissivity paint, effectively creating 14 unique target types. Each of these 14 was rotated at 2 to 3 different speeds and at 2 to 3 different angles (the axis of rotation pointing approximately orthogonal to the sensors (orthogonal spin), pointing approximately 40 degrees to one side (both orthogonal and parallel components to spin), and pointed straight toward the sensors (parallel spin)). There was more or less "wobble" of the targets through each of their revolutions as well as small drift in their spin speed and axis of rotation over the course of the 5-second data-gathering period that corresponded to a single run of the experiment.

B. Targets

Below are listed the 7 separate shapes used for targets in both unpainted and painted configurations. The following pictures show the majority of the targets, unpainted, with completely nonmetallic connectors and hooks (originally for attaching to the zip line):

A. 12" flat thin steel Square

B. 12" bent thin steel Square (final dimensions 12" x 10" x 3.25")

C. Small Cylinder (11" long x 8" diameter)

D. Medium Cylinder (16" long x 12" diameter)

E. Long Cylinder (5' long x 8" diameter)

Figure 1. Small Cylinder Target.

F. "Bullet-Shaped" (Actually a 12'' hemisphere jointed to a 12'' x 12'' cylinder)

G. A complex shape (designed to be overall similar to the Bullet-shaped Target and the same overall dimensions but with 4 extra metal objects (a 12" x 12" flat plate, a small closed cylinder, and two open rectangles) joined to it and open at the cylinder end instead of closed). We suspect that this object would look like the bullet in IR but different from it in the radar domain.

Figure 2. Medium Cylinder Target.

For each stationary rotation experiment, the target was attached, the fan regulator voltage set, the target started spinning, and then both IR cameras and the radar were started approximately simultaneously. When the 5-second data-gathering period was complete, the fan was turned off. Initial experiments were at night. We determined that the LWIR camera gave good results even in the daytime, so we decided to perform the majority of the experiments during the day to be able to also consider the realistic situation when the sun would illuminate real targets in space.

Figure 3. Long Cylinder Target.

Figure 4. Bullet-Shaped Target.

Figure 5. Complex Shape.

Figure 6. Complex Shape, Different Angle.

C. Data Processing

IR Low-Level Processing

The sky (which happened to be cloudy in all experiments) was used as the background for all IR collected data. Each IR frame was independently down-sampled from the over-resolved image. Down-sampling was done to simulate imagery more closely representative of real operating conditions. The down-sampling was sufficient to reduce the target to less than one pixel. Note that this involved both randomizing the location of the object within the

macropixel independent of either its actual location or its location in the last frame (to simulate an object moving at very high speed skipping many pixels from one frame to the next) and averaging the irradiance of all the actual pixels in each macropixel. The macropixels were square where each dimension was at least twice the size of the longest dimension of the target across all frames, thus ensuring that the object was always smaller than half a pixel. We termed these macropixels, "50%-macropixels" because the object was always guaranteed to be less than 50% of their size. We also did some rotational frequency-finding experiments with 40% and 20%-macropixels with very good results. (The rotational frequencies found from processing 20%- and 40%-macropixels were all either identical or very close to the ones found processing the 50%-macropixels.) The final step was to segment the target from the background. This was done by selecting the macropixel with the greatest irradiance and then writing this pixel to a file along with the 8 surrounding pixels (to capture the situation that the object may have crossed a macropixel boundary). Thus each frame of IR data was reduced to a total of 9 gray-scale values.

Radar Low-Level Processing

The data matrix was processed via the short-time Fourier transform using a Hamming window. The window size plays an important role in determining the quality of the image: a small window is able to quickly detect changes in frequency domain, but the integration gain is low, leading to higher noise and signal artifacts. Conversely, a large window produces a high SNR and smoother dynamics, but the long dwell time may cause us to miss small Doppler variations that are important for feature extraction.

In this preliminary effort, the Hamming window size was chosen heuristically to be 128. Through signal processing, the optimal window size would be determined in future work by solving a constrained minimization problem that maximizes SNR while allowing small features (changes) to be detectable.

Example range Doppler maps at maximum and minimum energy return points in the cycle are shown. (It should be noted that these data were not filtered to remove noise, did not have the background clutter removed, was not normalized, did not have its window length or function optimized to give the best results, and did not have its power adjusted. In fact, the radar has not even been calibrated, owing to delays in receiving parts and having the radar assembly completed immediately before the experiments began.)

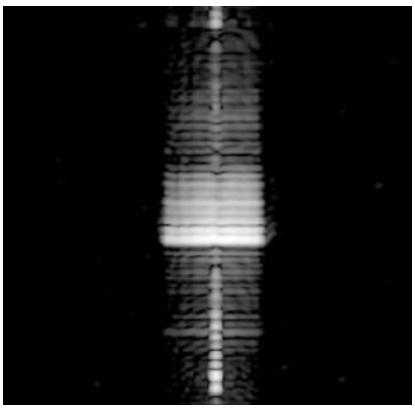

Figure 7. Radar Range Doppler Map at Maximum.

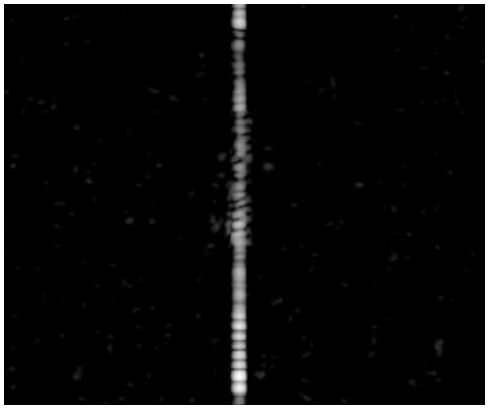

Figure 8. Range Doppler Map at Min.

High-Level Feature Extraction

Several high-level features were extracted from both the IR and the Radar signals, the most prominent of which was the rotational frequency of the target. The technique for doing this has been surprisingly simple. For the IR data, the average of the nine pixels was calculated and a cosine curve automatically best-fit matched to that average over time. This was done by finding the closest matching cosine as described in the Best-fit Cosine-Matching subsection further below. Similarly, for radar, the intensity of all the pixels on the range Doppler map was averaged and a cosine curve automatically best-fit matched to that average over time. We believe that the simplicity of this technique is responsible for its robustness. In every case, the actual rotation frequency was found to be half the frequency of LWIR signal best-fit cosine match, as expected. The best-fit cosine match of the average radar intensity was always found to be a multiple of 1, 2, or 4 of the actual rotational frequency. The agreement between radar-derived object rotational frequency and IR-derived rotational frequency had an average error of 1.4% and much of this was likely due to issues associated with low-cost experimentation, such as the spin vector of the rotation object actually drifting in speed or direction during the run. It should moreover be noted that these results were obtained at

sea level with objects only very slightly different from one another in temperature and occupying only a fraction of a pixel. Typically the amplitudes of the found best-fit cosine functions were 3 orders of magnitude smaller than the average irradiance. And the amplitudes were typically 4 orders of magnitude smaller in the case of the 20% macropixels. Yet the correct rotational frequency was always found due to the fact that the data were so consistent.

Other features were also extracted from the IR and radar data. Pixel crossings were counted and their frequency used to estimate dimensions. The error range, using a 95% confidence, was calculated to be 1/Sqrt(N) where N was the number of samples. Since many of our samples had 200 frames, this meant an error range of +/- 1/Sqrt(200) = 7%. There were 3 basic kinds of pixel crossings: horizontally across 2 pixels, vertically across 2 pixels, or across more than 2 pixels, which was termed "both." (Since all our objects were guaranteed to be smaller than a pixel, if more than 2 pixels were involved it had to be either 3 pixels in an "L" configuration (where the "L" could be oriented in one of 4 ways) or 4 pixels in a square. For the purposes of calculating the percentage of vertical crossing events, both "Vertical" and "Both" were counted. Similar counting was done for horizontal pixel-crossing events. A couple of example calculations are shown below:

E.g., Test #13 LWIR (Small square, 0 degrees orientation (spin vector orthogonal to sensor-to-target line), Macro pixel size was 0.6m)
200 samples – 101 cross horizontally or 50.5% (+/- 1/Sqrt(200) = 7%) so the average horizontal dimension is 0.505 x 0.6 m = .3m (compared to 0.3m actual)

- 63 cross vertically or 31.5% (+/- 7%) so the average vertical dimension is 0.315 x 0.6 m = .189m (compared to a range of 0 (edge on) to 0.3m (face on) actual) (the average of cosine is 0.636 and 0.189/0.3 = 0.63)).

E.g., Test # 18 LWIR (Small cylinder, 0 degrees orientation, Macro pixel size is 0.7m)
200 samples – 78 cross horizontally or 39% (+/-7%) so the average horizontal dimension is 0.39 x 0.7 m = .27m (compared to 0.23m actual if the orientation was perfect)

– 84 cross vertically or 42% (+/- 7%) so the average vertical dimension is 0.42 x 0.7 m = .29m (compared to 0.2 (edge on) and 0.28m (face on) and 0.34 tipped up to maximum vertical size (actual average over time is approximately 0.29m)) .

These features were extracted automatically by the prototype and the maximum of the two was used as Size for correlation with radar-derived size.

Another set of features that can be extracted relates the minimum and maximum irradiances and cross-sectional areas. For example: in test run #26 a minimum occurs at 0.133 seconds of 1,113,130,000 and a maximum at 0.333 of 1,116,660,000.

This might seem like a very small difference. However, the empty sky in that frame had an average irradiance of 1,111,250,000 so the object in the minimum frame contributed 1,113,130,000 - 1,111,250,000 = 1,880,000 and in the maximum frame contributed 1,116,660,000- 1,111,250,000 = 5,410,000 so the ratio of the maximum contribution over the minimum contribution is 541/188 = 2.9. The ratio of the maximum area to the minimum area presented by the object to the camera is thus approximately 3. In this particular example, this object is a large but empty cylinder and the minimum occurs during its rotation when the camera happens to site down its opening.

Size was also extracted from the radar data. As described above, frames containing the range-Doppler map were output by low-level radar processing and all of the pixels of every frame averaged to find the frames with the maximum and minimum returned energy. For frames with maximum returned energy, simple rectilinear edge detection was performed. (I.e., bounding boxes with edges parallel to the range and Doppler axes were found.) The resulting pixel locations of the object's minimum and maximum Doppler signal were first converted into Hz then into a velocity giving both a Left and Right velocity number. These could be converted into distances with the measured rotational frequency. Adding them together gave the radar-derived Size measure used by the prototype for correlation. Their ratio could also be used as a measure to the degree that the center of gravity (the center of the spin axis) was offset from the center of the object's width. This could be used for both correlation and classification, especially when it was far off center. Some example calculations are shown below.

Radar max rotation width, cg unbalance
E.g., Test #13 Radar (Small Square, 0 degrees orientation)
Left Doppler Edge: pixel 166 -> -117.51 Hz -> -2.150 m/second
Width = V/(2pi x rotationalFreq: 2.2 (half LWIR signal freq)) = -0.155m
Right Doppler Edge: pixel 262 -> 103.69 Hz -> 1.897 m/s -> 0.137m
Total rotational Width = 0.155 + 0.137 = 0.29m (versus 0.30 m actual)

Off Center: (0.155 – 0.137)/2(0.155 + 0.137) = 3%

E.g., Test #18 Radar (Small Cylinder, 0 degrees orientation)
Left Doppler Edge: pixel 181 -> -82.95 Hz -> -1.517 m/second
Width = V/(2pi rotationalFreq: 1.60 (half LWIR signal freq) -> -0.151m
Right Doppler Edge: pixel 260 -> 99.08 Hz -> 1.812 m/s -> 0.180m
Total rotational Width = 0.33m (versus 0.34 m actual)
Off Center: (0.151 – 0.180)/2(0.151 + 0.180) = 4%

IV. Results

A. Radar/IR Experiment Results

The MW and LW video could be easily observed in real time while the data were being collected. The rotations could be clearly seen in the high-resolution video. It did appear that for all the targets that the total irradiance of the target was oscillating at twice the frequency that the target was rotating. The unpainted targets primarily were reflecting the radiation from the sky and ground (so that the maximum irradiance from the flat plate, for example, tended to be when it was at 45 degrees and oriented to reflect the ground and the minimum tended to occur when it was oriented at 45 degrees, reflecting the sky). The black-painted targets tended to have the maximum and minimums based on the cross-sectional areas they presented to the camera. The radar signals could not be observed in real time, but with a few minutes of processing, AVI files could be created that appeared to show the total returned energy oscillating with the same frequency as the IR signal. Some crude timing calculations were done using individual frames to calculate the period of oscillation of the object and periods of oscillation of the IR and radar signals. The LW camera was shown to be largely insensitive to sunlight, while the MW signal often had sunlight glint.

Experimental data were analyzed and processed for 40 runs. Several runs were not analyzed because they were essentially redundant and due to the scope of a low-cost effort. The 40 runs included all 7 shapes, both painted and unpainted (essentially 14 unique targets) for a variety of spin speeds and 3 very rough directions. The directions were approximately rotating straight toward and away from the sensors (called 0 degrees), the spin vector pointing approximately right at the sensor so the rotational movement occurred across the sensor's view (called 90 degrees) (and which we assumed would not allow rotational frequency to be discerned with our implemented techniques), and somewhere in between the 0-degree and 90-degree cases, slightly closer to the 0-degree case (called 40 degrees). This third state was intended to provide a middle ground but not be precisely in between the other two.

The 40 rows of analyzed data are shown below. Almost all runs have both LWIR and radar data. A few have Midwave IR (MWIR) data. As mentioned above, MWIR was more sensitive to sunlight and so was not normally used for automatic analysis. (It did often prove useful to establish ground truth since it had a higher and more consistent frame rate than LWIR. In a couple of cases LWIR was not recorded or was otherwise unusable). As can be seen from the table, there is remarkable agreement between radar-derived and the IR-derived rotational frequency of the objects. (The table actually lists twice each object's rotational frequency, essentially the LWIR view.) The average error is only 1.4% and is only above 4% in one case and that was for the most rapidly rotating object recorded when the LWIR's limited frame rate started to become an issue. Much of this 1.4% average error was attributable to low-cost experiment factors that would not be present during operational use such as the fact that there was actual variation in the rotational frequency and direction of the object during the data gathering and the time between frames of the LWIR camera varied considerably and unpredictably.

The error is also consistently low across all objects, whether they were painted with high- emissivity paint or not. This is very interesting because the techniques were able to function very effectively even when the IR energy received by the camera from the target was primarily reflected. The error is also low for both the 0-degree and 40-degree cases, showing that the techniques are not dependent on rotation being "lined up" with the sensor. Surprisingly, there were good results for even some of the 90-degree cases, though the error tended to be higher. These are probably a result of the spin vectors not being precisely lined up for the whole time period of the data runs. It does illustrate how robust the techniques are. The error was also low across a wide breadth of rotational frequencies—from about 0.45Hz to about 3.2 Hz object rotational frequency. The techniques would easily work on slower-rotating objects as long as data were gathered for at least 2 periods, preferably 3. (For simple, symmetrically shaped objects, 1 period would be enough, since the signal would repeat twice in one rotational period.) On the upper end, a faster frame-rate camera would be required for faster-rotating objects.

The low error percentage is especially impressive given that no calibration was done for either the radar or IR cameras. This is contrary to the philosophy of many in the missile defense sensor field where very large efforts are thought to be mandatory to collect large quantities of data to calibrate and seed databases. While this effort is certainly useful and probably very valuable, it is not necessarily required for all possible techniques— as this effort illustrates. (And certainly calibration and other data gathering would allow additional features to be extracted and utilized here as well.)

Test #	Target	Paint	Approx Initial Angle	MWIR Freq	LWIR Freq	Radar Freq	error
13	Small Square	No	0		4.337	4.424	1.97%
17	Small Square	No	40		4.417	4.266	3.42%
18	Small Cylinder	No	0		3.201	3.235	1.05%
20	Small Cylinder	No	0		4.114	4.098	0.96%
23	Small Cylinder	No	40		4.05	4.019	0.77%
26	Big Cylinder	No	0		2.59	2.59	0.00%
31	Big Cylinder	No	40		3.344	3.36	0.48%
32	Complex	No	0		3.354	3.366	0.36%
34	Complex	No	0		4.608	4.639	0.67%
35	Complex	No	0		2.718	2.782	2.30%
39	Complex	Yes	0		2.718	2.722	0.15%
43	Complex	Yes	40		2.033	2.049	0.78%
47	Small Square	Yes	0	4.64	4.623	4.544	1.71%
49	Small Square	Yes	0	4.704	4.655	4.672	0.36%
51	Small Square	Yes	40	4.926	4.91	4.8	2.24%
53	Small Square	Yes	90		2.145	2.145	0.00%
54	Small Square	Yes	90		2.59	2.574	0.62%
55	Small Cylinder	Yes	0		4.162	4.146	0.38%
56	Small Cylinder	Yes	0		6.927	6.621	4.42%
57	Small Cylinder	Yes	40		6.434	6.322	1.74%
58	Small Cylinder	Yes	40		3.631	3.653	0.60%
61	Big Cylinder	Yes	0	2.07		2.081	0.53%
62	Big Cylinder	Yes	0		2.936	2.92	0.54%
63	Big Cylinder	Yes	40		4.162	4.114	1.15%
64	Big Cylinder	Yes	40		2.033	2.017	0.79%
65	Long Cylinder	No	0		1.072	1.072	0.00%
66	Long Cylinder	No	40		0.955	0.939	1.68%
67	Long Cylinder	No	90		0.462	0.446	3.46%
68	Long Cylinder	Yes	0	0.939	0.923	0.923	0.00%
72	Long Cylinder	Yes	40		1.9	1.852	2.53%
74	Warhead	No	0		4.066	4.05	0.39%
75	Warhead	No	0		3.694	3.64	1.46%
76	Warhead	No	40		2.415	2.431	0.66%
78	Warhead	No	90		2.288	2.208	3.50%
81	Warhead	Yes	0		3.917	3.981	1.61%
87	Curved Plate	No	0		3.312	3.446	3.89%
88	Curved Plate	No	0	4.035		4.082	1.15%
89	Curved Plate	No	0	5.473	5.473	5.346	2.32%

Figure 9. Experiment Results

B. Prototype Description

Prototype Features Used

Normal expected target spin rates will usually be significantly less than the frame rate of the IR sensor. For the component of spin orthogonal to the IR sensor, if the object is not approximately symmetric to the spin axis, the visual cross section (and therefore the intensity) will fluctuate. The frequency of this fluctuation will often be twice the spin frequency (since most objects have 2 maximum cross section rotational angles, i.e., two sides) but may be 4 or 6 times the spin frequency (for square and triangular shapes, though these would involve more subtle fluctuations). Similarly, the RCS will fluctuate in a similar fashion with a frequency related to the orthogonal spin frequency, though the frequency variation may be more complex and have a single peak per rotation (instead of two). A micro Doppler radar signal will also show the speed of the end points. If the radial width fluctuates through the revolution, this fluctuation will typically be at twice orthogonal spin frequency. The approximate maximum radial length can be estimated from the Doppler maximum indication of speeds and measured orthogonal angular rotation rate (spin frequency). Using the calculations described above, the prototype also extracted and used the object's size from both the IR and radar signal.

Likelihood Calculation that 2 Sensed Targets are the Same Object

For a pair of sensor observations that need to be fused (i.e., one from radar and one from IR), for every possible (i.e., the metric tracking data and covariances are such that the two targets could be the same object) pairing of a target in one set to a target in the other set, the software needs to calculate the probability that the two targets are the same object. This is basically a matching and probability calculation. For each of the features extracted from the sensor data for one sensor type, one must determine the degree that it matches the same feature extracted from the other feature type, given that the other sensor observes the object from a different angle and distance. This degree of matching, when the feature exists in both domains, is primarily based on the degree to which their error ranges overlap. In the case where the feature was not able to be estimated from the other sensor's data, the software must determine the probability of this occurring, given the assumption that the targets are the same object. For example, if the first sensor is IR and it determines a specific frequency of rotation for the target, and the second sensor is radar and no rotational frequency is evidenced, but the viewing angle for the radar is such that the spin axis direction could be parallel to the radar's view (consistent with the estimated parallel and orthogonal components from the IR data), then this factor must be considered. But if the angles, sizes, and frequencies were such that the radar really should have discerned the frequency, then it should reduce the probability of a match. Some of the features can be extracted in multiple ways, even in just one sensor domain. In these cases, if the feature is contradictory, then matching either value to the same feature in the other object should increase the probability that they are the same. (For example if two methods executed on the IR data predict that the spin frequency is 1Hz and 3.5 Hz and the spin frequency extracted from the radar data is 3.5 Hz, then that should be considered a perfect match.)

The matches, lack of matches, and mismatches for the pair of objects must be combined into an overall probability estimate for associating the pair. A common way to estimate probabilities based on evidence is to use Bayesian Networks. This does require a priori and dependent probabilities to be calculated. In future work, more rigorous or statistical methods can be used, if enough data exist. Case-based reasoning (CBR) is another possibility if enough measurements have been performed for there to be sufficient cases to support it. There are only M x N such calculations to consider, where M and N are the number of targets in each set, so this is not computationally burdensome.

Based on the extracted features for each object, the probability that the object was not sensed in the other domain should be calculated as well. For example, an object known to be small from its RCS is less likely to be sensed by IR than a large one. Since the sensitivity and range of the IR sensor are known, and a range of reasonable temperatures for the object can be assumed, the probability of detection across that range for the object for the specific IR sensor (given its distance) can be calculated. There are only M + N of these calculations.

Determine Most Likely Consistent Association Sets (Hypothetical Associations Sets)

Given the calculated probabilities (in isolation) that any pair of objects match and the probabilities that an object was undetectable in the other domain, maximally consistent sets of associations can be generated. The main problem, of course, is that even the most likely associations may not all be true. And, later, additional evidence may become available that will confirm or refute various hypothetical associations. The solution is to use multiple-hypothesis reasoning to keep track of different hypothetical worlds of assumptions and the dependencies between them.

Of course, the first step is to construct the most probable consistent sets of associations. Although there are only (at most) (M x N) + (M + N) pair wise probability calculations (including the undetected cases), there is an exponential number of possible association sets. But given that the orbit metric tracking data and covariances limit the possibilities, this number may still be manageable. An algorithm that works well in both small and large number cases is to first make an ordered list of the (M x N) + (M + N) possible associations sorted in decreasing probability. One should construct the first set by starting at the top of the list, working down, picking the first association where both targets have not been chosen yet. Depending on how many associations with undetected objects are selected, this set will be at least Max(M,N) and at most M + N long. If M and N are relatively small, it would be possible to systematically assemble every possible combination, calculate its total probability, and keep the most likely ones. If this is not practical, a good heuristic for assembling the second and subsequent ones is similar to assembling the first set except that the first selection (instead of just being from the top of the list) should be one of the likely associations that does not appear in any list yet. After enough sets have been generated to exhaust this first selection method, then the first two choices can be replaced (instead of just the first one) by the most likely consistent pairs of associations that have not appeared together yet. Depending on the size of M and N, this process can proceed to triples and beyond, until enough hypothetical sets are generated that their cumulative probability is very high. (I.e., it is unlikely that the true set of associations is not represented by one of the hypothetical association sets.)

Within Hypothetical World Propagate Each Object Forward to Additional Sensing Opportunities

Many orbital (and sub-orbital) propagation programs exist that can, given orbit metric data, determine the location of the object at a specific future point in time. A few will also propagate covariances forward. Particle techniques can also be utilized to use propagators to propagate covariances. For each hypothetical set, the associated targets can be assumed to be the same object and their covariances combined, which, in the case of fusing radar and IR observations, is usually especially beneficial, since orbit metrics from IR and radar tend to be very complementary (radar being most accurate in range and IR being most accurate in angle). Using the combined (reduced) covariance and orbit metric data, the hypothetically fused object can be propagated forward (using one of the many existing propagators) to the next sensing opportunity and the probability of detection of the object by that sensor calculated based on the fused object's features (size, temperature, rcs, etc.). (Of course, since different hypothetical sets share associations, the combined covariances and forward propagated locations should be cached and reused.)

The follow-on sensing opportunities fulfill two purposes. First they tend to confirm (especially if the combined covariance is significantly smaller than either alone) or refute (if the probability of detection is reasonably high) the associations and therefore confirm or refute the hypothetical association sets that the confirmed or refuted association is a part of. Second, in the cases where the association of the 2 targets from the first two sensing opportunities is confirmed by a closely matching object in the third sensing opportunity, they extend the associations and hypothetical sets.

The third (and subsequent) sensing opportunity is matched with fused objects in each hypothetical set in exactly the same way as described above with the exception that more data exist for the hypothetically fused objects because they include data from different types of sensors and different opportunities, all combined into one hypothetical object. In particular, features may exist from a previous sensing opportunity for the same type of sensor (though at a different time and from a different perspective). Over time, the associations in the hypothetical association sets grow from the initial pairs of targets to triples and then to quadruples and so on. As specific associations are confirmed/refuted (their calculated probability grows/shrinks), the hypothetical association sets they are part of are retained/pruned away.

Best-fit Cosine Matching

The most obvious way to determine the spin frequency is to use the FFT or DFT and examine the most prominent component. However, for a small noisy sample, there may be some significant additional frequency components introduced. So an additional, very noise-tolerant method for determining the rotational frequency was included—best-fit cosine matching.

This module searches the space of cosine functions to find the cosine parameters that most closely match an observed data set. Match accuracy is determined by a least-squares comparison between the hypothetical cosine function and the total data set. Parameters calculated by this procedure are:

- Frequency.
- Amplitude.
- Phase.

The matching process is carried out by a brute-force search of potential cosine parameters within a given range. While already relatively fast in its current implementation, the algorithm is highly parallelizable, and it could easily be made more efficient by intelligent searching versus the brute-force approach. Examples of cosine finding are shown below for both radar and LWIR for test # 49. In the figures below, the top part of the figure is the total returned energy of the signal and the bottom part is the closest matching cosine wave. Irregularities in the matching cosine wave are a result of sampling.

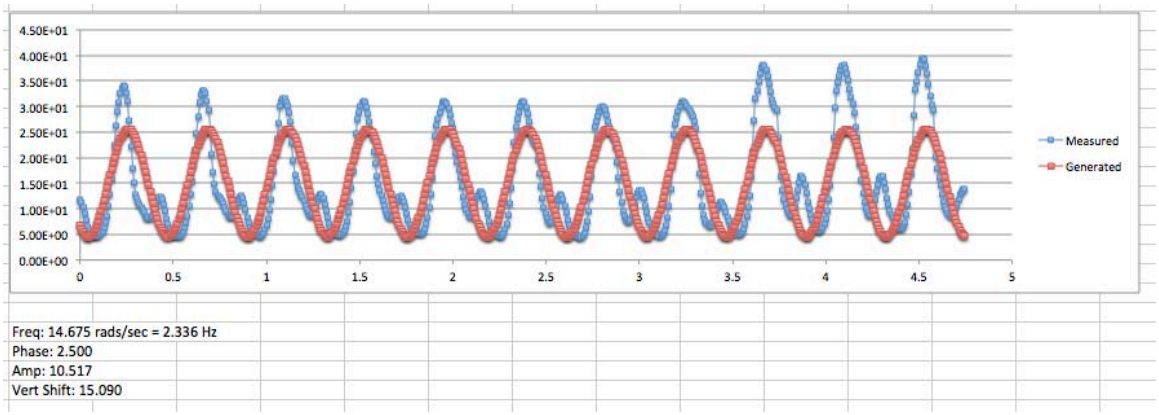

Figure 10. Cosine found for Test # 49 radar.

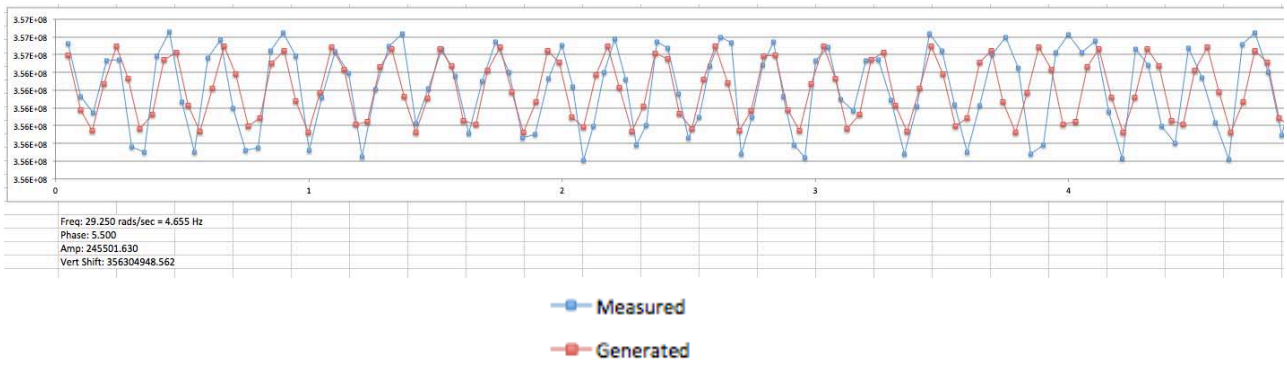

Figure 11. Cosine found for Test # 49 LWIR.

Hypothetical World Generation

We implemented in the prototype an Object-Oriented architecture and supporting code to support the generation, maintenance, and tracking of a large number of hypothetical "world states" representing possible configurations of objects from different sensor arrays at different times. Every time a new observation set is introduced to the system, all possible resultant object association sets are generated from the data in that observation.

Feature Correlation

In order to support the calculation of similarity between objects in different observation sets, we implemented in the prototype a modular approach to feature correlation. A set of Feature Correlator objects encapsulate the algorithms used to determine similarity between identical attributes from identical or different observation types (e.g., Radar vs. IR). By abstracting the algorithmic implementations of these calculations into their own software objects, our system supports the addition of new attribute types or different algorithmic approaches without requiring a re-architecting of the underlying representations of objects, observations, and associations. Examples of feature types include:

- Size/Shape/Dimensions.
- Rotation Frequency.
- Velocity.

Generally, each feature extracted for an object from the sensor data has an estimated range of possible values. For example, rotational frequency tends to have a narrow range even for our low-cost experiments (when there was a lot of clutter and noise), but radar-generated size tended to have a wide range. Matching of the same feature for different observations of potentially the same object has to take these ranges into account, generally assuming the probability relates to the degree of overlapping. Additionally, for the rotational frequency and the radar-calculated size (which uses rotational frequency), the prototype supports harmonic matching. E.g., an IR-derived frequency could either match the radar-derived rotational frequency or a small integer multiple of it.

C. Scenario Execution Results – 9 different ones, 3 are presented here

Nine separate scenarios were constructed from the experiment runs. Each scenario consisted of a set of observations and each observation referenced a set of IR or radar data files recorded from our expereiments. The set of scenarios ranged from simple examples (two scenes of two objects each) to ones that were more complex, including missing objects in some scenes, having more than 2 scenes, and including up to 8 objects. The prototype successfully handled all situations presented to it. Each scenario is briefly listed below along with the the final prototype's estimated correlations.

Scenario 2 – A simple Example with a missing object
Observation 0: LWIR
- Object 0: 049 (Small Square)
- Object 1:075 (Warhead)

Observation 1: Radar
- Object 2: 049 (Small Square)
- Object 3: 075 (Warhead)
- Object 4: 018 (Small Cylinder)

Results: 4 post-pruned worlds (Objects grouped together were determined to be the same object):
[4 (1, 3) (0, 2)] Probability: 0.22
[1 4 3 (0, 2)] Probability: 0.033
[0 4 (1, 3) 2] Probability: 0.026
[1 0 4 3 2] Probability: 0.0039

Scenario 7 - Large
Observation 1: Radar
- Object 0: 013 (Small Square)
- Object 1: 020 (Small Cylinder)
- Object 2: 097 (Curved Plate)
- Object 3: 077 (warhead)
- Object 4: 065 (Long Cylinder)
- Object 5: 026 (Big Cylinder)
- Object 6: 023 (Small Cylinder)
- Object 7: 049 (Small Square)

Observation 2: LWIR
- Object 8: 013 (Small Square)
- Object 9: 020 (Small Cylinder)
- Object 10: 097 (Curved Plate)
- Object 11: 077 (warhead)
- Object 12: 065 (Long Cylinder)
- Object 13: 026 (Big Cylinder)
- Object 14: 023 (Small Cylinder)
- Object 15: 049 (Small Square)

Results: 12 post-pruned worlds:
[(3, 8) (4, 12) (2, 10) (1, 9) (6, 14) (5, 13) (7, 15) (0, 11)]Probability: 0.074
[(0, 8) (4, 12) (2, 10) (1, 9) (6, 14) (5, 13) (7, 15) (3, 11)] Probability: 0.066
[(3, 8) (4, 12) (2, 10) (7, 9) (6, 14) (5, 13) (1, 15) (0, 11)]Probability: 0.048
[(0, 8) (4, 12) (2, 10) (7, 9) (6, 14) (5, 13) (1, 15) (3, 11)] Probability: 0.042
Probability: 0.026
Probability: 0.024
Probability: 0.017
Probability: 0.016
Probability: 0.015
Probability: 0.013
Probability: 0.011
Probability: 0.010

Scenario 8 – 3 Observations, each one missing a different object
Observation 1: LWIR
- Object 0: 064 (Big Cylinder)
- Object 1: 068 (Long Cylinder)
- Object 2: 089 (Curved Plate)

Observation 2: radar
- Object 3: 064 (Big Cylinder)
- Object 4: 088 (Curved Plate)
- Object 5: 089 (Curved Plate)

Observation 3: MWIR
- Object 6: 068 (Long Cylinder)
- Object 7: 088 (Curved Plate)
- Object 8: 089 (Curved Plate)

14 post-pruned worlds:
[(0, 4) (3, 7) (2, 5, 8) (1, 6)] Probability: 0.00283
[4 (0, 3, 7) (2, 5, 8) (1, 6)] Probability: 0.00263
[3 (0, 4, 7) (2, 5, 8) (1, 6)] Probability: 0.00057
[(0, 3) (4, 7) (2, 5, 8) (1, 6)] Probability: 0.00047
[(2, 5) (0, 4) (3, 7) 8 (1, 6)] Probability: 0.00041
[(2, 5) (0, 4) (3, 7) 8 (1, 6)] Probability: 0.00041
Probability: 0.00038
Probability: 0.00033
Probability: 0.00036
Probability: 0.00016
Probability: 0.00014
Probability: 0.00012
Probability: 0.00011
Probability: 0.00011

V. Future Work

There are many possible activities that could follow on from these efforts. Initially, we necessarily focused on the reasoning unique to the sensor signals we had collected and the features that we were able to extract within this limited scope and not on all techniques that might also be useful in an operational system. Further research could

also be performed as to different types of inferencing that could be made based on an expanded set of extracted features. This includes more complicated 3-dimensional reasoning relating to geometric inference based on pose and the constraints on dimensions that each sensing opportunity provides.

For a specific application, the details of the specific sensors and the specific objects and shapes of interest could be more thoroughly utilized to determine what kind of data are available to what level of resolution in the sensor signals. Combining this with an understanding of what types of scenarios should be expected (e.g., multiple objects in the same radar beam (or less likely, one IR pixel), maneuvering and deploying objects, possibly complex shapes, possibly more complex spin situations, expected distribution of object spacing, etc.) would provide better guidance as to which lines or automatic reasoning should be pursued. For example, we primarily used macropixels sized to ensure that objects were smaller than 50% of the macropixel size and only performed a few experiments to confirm that the 20% cases also worked well. But real knowledge of the target sizes and IR resolution would provide guidance as to what ratio to pursue most vigorously. Given that multiple objects may be in the same radar beam and reasoning will be required to extract the separate frequency components, it is fortunate that the radar signal tends to exhibit a purer repeating pattern.

Another issue for more future work is reasoning with metric data. In this low-cost effort, the scenarios all assumed that any object in one scene could be any object in another scene for purposes of correlation. But in an actual operational system an object's sensed position, velocity, and error covariances limit which other objects it could possibly be. (We did implement this capability in this prototype but there was no plausible way to realistically demonstrate it from the collected data.) Similarly, there was no way to calculate a probability of an object being detected (or not detected) in another scene and so the prototype just used the same constant for all cases. But an operational system would be made more accurate and efficient by using the features extracted from one scene to determine the likelihood that the object will appear in another sensor's scene. In this effort, no temperature information was extracted from the IR signals. Since this is important information for classification, this would need to be performed in future work. Similarly, since the radar system was not calibrated, we could not determine radar cross section from our radar returns or compare the radar returns from separate runs (i.e., to determine the relative sizes of two different objects). Obviously this would be remedied in future work.

The cosine finder was a very robust mechanism for determining rotational frequency from both radar and IR data but other possibilities should be investigated. Fourier transforms (Discrete Time Fourier Transform (DTFT)) were investigated on a sample of data and were found to not be as robust as the cosine finder and not as unambiguous. (They correctly found the rotational frequency about 80% of the time). But this does not mean that they would not have useful information to offer. Similarly, many of the signals were very consistently repeating but with very unusual, strange looking patterns. This initiated the concept of a repeat pattern finder, which would quickly search for an odd-shaped repeating pattern by quickly trying different time spans and seeing how closely those time spans repeated. With multiple ways to determine rotational frequency, their results could be combined in various ways—looking at when they agree and disagree for example and using this as a confidence in the result. Similarly, the results of DTFT could be used to "seed" the pattern repeat finder or the cosine finder. Further research is needed to determine the relative benefits of different combinations of approaches on actual returned IR and radar signals.

More accurate experiments could be done in terms of more accurately manufactured targets, more steady and driftless rotations (speed and direction), more accurate determination of relative angles of sensors, targets, and spin vectors, more variations in speed, angle, and axis of rotation, and better elimination of spurious signals and clutter. These would help determine the accuracy limitations, which we believe to be very high, for sensed rotational frequency. In this effort, the radar environment was much more cluttered and noisy than would be expected during actual missile defense operations. Therefore in the future, we plan to use an anechoic chamber to gather radar data on a variety of objects of interest for missile defense, with a variety of spin vectors, from a variety of sensor angles, and with various polarizations.

In the next stage of development, another possibility is to investigate a microscan implementation in which the point spread function produced by the optics is adjusted to be greater than the pixel size (potentially possible with recent advances in focal plane technology). In this case, the imagery is properly sampled and hence a variety of deconvolution and processing techniques can be applied. Microscanning, sometimes called "super resolution," (either deterministic or random based on vibration) can allow for increased sampling beyond the pixel pitch of the detector. Such an implementation is possible because of advances in processing technology. This investigation would define what size/shape information of the target may be accessible using these techniques.

A variant of microscanning relates to the integration time of the IR camera and the high speed of the objects. From this combination, one would expect that a single object in a single frame would actually be detected as a "streak" of several pixels, crossing several pixel boundaries in the direction of travel. Since these streaks would likely cross several horizontal and vertical pixel boundaries, there is an opportunity to derive some subpixel size

information from even a single frame, thus increasing the power of the pixel crossing techniques previously discussed.

VI. Conclusions

We can state definitively that it is possible to extract high-level object features (such as rotational frequency, size, and dimension ratios) from the low-level IR and radar signals, even when an object is less than the size of a single IR pixel or radar beam. We can also state that these features can be used to correlate the same object in different scenes sensed by dissimilar sensors. We were able to demonstrate our techniques with real radar and IR data gathered from real objects with a very high precision (low average error percentage). There appears to be enormous potential for this work to have a major impact on correlation and classification of objects in space.

These results were accomplished without much computational processing power. Because the objects are smaller than one pixel, the IR calculations only involved at most 9 pixels per frame (the pixel the object was definitely in and all 8 immediately surrounding pixels). Extracting low-level features only involved simple math operations like summing, averaging, differences, and multiplications and so was essentially instantaneous. Similarly, processing the radar video frames was very straightforward. As input to feature extraction, we used the highly standard range-Doppler map and only performed simple operations such as averaging and finding the left and right vertical edge based on total column intensities. And these calculations would be even more efficient in an operational context owing to reduced noise, clutter, and lack of multipath effects. These would greatly reduce the effective size of the frame since it would only consist of a relatively narrow band in range (corresponding to the actual size in depth of the object).

And performing the correlation was also efficient since it used symbolic reasoning—not search or iterative optimization. The most time-consuming aspect of the prototype was maintaining a large number of hypothetical worlds, since every non-zero possibility was retained. But more efficient mechanisms for keeping the number within any limit required by computational bounds exist.

Perhaps most remarkable was how robust the object rotation frequency extraction techniques were. We believe this is because of their simplicity. Good results were obtained for all conditions tested, even without calibrating the sensors or the use of past recorded data of any kind.